Learning to Weight Samples for Dynamic Early-exiting Networks

Yizeng Han¹[∗][®], Yifan Pu¹[∗][®], Zihang Lai²[†][®], Chaofei Wang¹[®], Shiji Song¹[®], Junfeng Cao³[®], Wenhui Huang³[®], Chao Deng³[®], and Gao Huang¹^[⊠][®]

 ¹ Tsinghua University, Beijing 100084, China {hanyz18, pyf20, wangcf18}@mails.tsinghua.edu.cn {shijis, gaohuang}@tsinghua.edu.cn
 ² Carnegie Mellon University, Pennsylvania 15213, United States zihangl@andrew.cmu.edu
 ³ China Mobile Research Institute, Beijing 100084, China {caojunfeng, huangwenhui, dengchao}@chinamobile.com

Abstract. Early exiting is an effective paradigm for improving the inference efficiency of deep networks. By constructing classifiers with varying resource demands (the exits), such networks allow *easy* samples to be output at early exits, removing the need for executing deeper layers. While existing works mainly focus on the architectural design of multi-exit networks, the training strategies for such models are largely left unexplored. The current state-of-the-art models treat all samples the same during training. However, the *early-exiting* behavior during testing has been ignored, leading to a gap between training and testing. In this paper, we propose to bridge this gap by sample weighting. Intuitively, easy samples, which generally exit early in the network during inference, should contribute more to training early classifiers. The training of hard samples (mostly exit from deeper layers), however, should be emphasized by the late classifiers. Our work proposes to adopt a *weight prediction network* to weight the loss of different training samples at each exit. This weight prediction network and the backbone model are jointly optimized under a meta-learning framework with a novel optimization objective. By bringing the adaptive behavior during inference into the training phase, we show that the proposed weighting mechanism consistently improves the trade-off between classification accuracy and inference efficiency. Code is available at https://github.com/LeapLabTHU/L2W-DEN.

Keywords: Sample weighting, Dynamic early exiting, Meta-learning

1 Introduction

Although significant improvements have been achieved by deep neural networks in computer vision [25,39,40,13,18,7,33], the high computational cost of deep

^{*} Equal contribution. \boxtimes Corresponding author.

[†] Work done during an internship at Tsinghua University.

2 Y. Han, Y. Pu et al.



Fig. 1: **Our sample weighting strategy**. At test time, the classifiers at varying depths handle inputs with different complexity (top). However, the conventional training strategy (middle) treats all the samples equally at multiple exits. In contrast, our weighting mechanism (bottom) enables each classifier to emphasize the samples that it is responsible for.

models still prevents them from being applied on resource-constrained platforms, such as mobile phones and wearable devices. Improving the inference efficiency of deep learning has become a research trend. Popular solutions include lightweight architecture design [16,49], network pruning [27,28,34,47], weight quantization [20,10], and dynamic neural networks [11,36,17,46,31,1,45,43,12].

Dynamic networks have attracted considerable research interests due to their favorable efficiency and representation power [11]. In particular, they perform a data-dependent inference procedure, and different network components (e.g. layers [45] or channels [1]) could be conditionally skipped based on the input complexity. A typical adaptive inference approach is *early exiting* [17,46], which can be achieved by constructing a deep network with multiple intermediate classifiers (early exits). Once the prediction from an early exit satisfies a certain criterion (e.g. the classification *confidence* exceeds some threshold), the forward propagation is terminated, and the computation of deeper layers will be skipped. Compared to the conventional static models, such an *adaptive* inference mechanism can significantly improve the computational efficiency without sacrificing accuracy. When making predictions with shallow classifiers for canonical (*easy*) samples, a substantial amount of computation will be saved by skipping the calculation of deep layers. Moreover, the network is capable of handling the non-canonical (*hard*) inputs with deep exits (Fig. 1, *Inference Stage*).

Existing works on dynamic networks mainly focus on designing more advanced multi-exit architectures [17,46]. A naive training strategy is generally adopted by summing up the cross-entropy losses obtained from all classifiers. More importantly, the loss of both *easy* and *hard* samples contributes equally to the final optimization objective (Fig. 1, *Conventional Training*), regardless of where a sample may *actually exit*. However, these exits at varying depths have different capabilities, and they are responsible for recognizing samples of varying complexity at test time. Such an adaptive inference behavior has been neglected by the naive training paradigm adopted in previous works [17,29,46].

In this paper, we propose to bridge the gap between training and testing by imposing sample-wise weights on the loss of multiple exits. Our motivation is that every classifier is required to handle a *subset* of samples in the adaptive inference scenario. Specifically, the early classifiers only need to recognize some canonical inputs, while the deep classifiers are usually responsible for those non-canonical samples. Therefore, an ideal optimization objective should encourage each exit to emphasize different training samples by weighting their loss. In a nutshell, the challenge for sample weighting in multi-exit models is two-fold: 1) the exiting decisions are made during *inference*, and we have no prior knowledge of where a specific sample exits; 2) setting proper weights is a non-trivial problem.

To address these challenges, we propose to automatically *learn* appropriate weights by leveraging a *weight prediction network* (WPN). The WPN takes the training loss from all exits as input, producing the weights imposed on the samples at every exit. We jointly optimize the backbone model and the WPN in a *meta-learning* manner. A novel optimization objective is constructed to guide the meta-learning procedure. Precisely, we mimic the test-time early-exiting process to find where the samples will exit during inference. The meta objective for each classifier is then defined as the loss only on the samples that *actually* exit at this classifier. Compared to the conventional training strategy, our specialized meta-objective encourages every exit to impose proper weights on different samples for improved performance in dynamic early exiting (Fig. 1, *Our Training*).

We evaluate our method on image recognition tasks in two settings: a class balanced setting and a class imbalanced setting. The experiment results on CIFAR[24], ImageNet [5], and the long-tailed CIFAR [4] demonstrate that the proposed approach consistently improves the trade-off between accuracy and inference efficiency for state-of-the-art early-exiting networks.

2 Related works

Dynamic early-exiting networks. Early exiting is an effective dynamic inference paradigm, allowing *easy* samples to be output at intermediate classifiers. Existing works mainly focus on designing more advanced architectures. For instance, BranchyNet [41] attaches classifier heads at varying depths in an AlexNet [25]. An alternative option is cascading multiple CNNs (i.e. AlexNet [25], GoogleNet [40] and ResNet [13]) to perform early exiting [2]. It is observed [17] that the classifiers may interfere with each other in these chain-structured or cascaded models. This issue is alleviated via dense connection and multi-scale structure [17]. Resolution adaptive network (RANet) [46] further performs early exiting by conditionally activating high-resolution feature representations.

4 Y. Han, Y. Pu et al.

While considerable efforts have been made on the architecture design, the *training* of these models still follows the routines developed for static networks. The *multi-exit* structural property and the *early-exiting* behavior are usually ignored. Some training techniques are studied to boost the knowledge transfer among exits, yet still neglecting the *adaptive inference* paradigm [29]. In this paper, we put forth a novel optimization *objective*, encouraging each exit to focus on the samples that they would probably handle in dynamic inference. **Sample weighting**. Different training samples have unequal importance. The idea of sample weighting could be traced back to dataset resampling [3] or instance reweighting [48]. These traditional approaches evaluate the sample importance with some prior knowledge. Recent works manage to establish a loss-weight mapping function [9,35,32,22,26,21]. There are mainly two directions: one tends to impose larger weights on *hard* samples to deal with the data imbalance problem (e.g. focal loss [32]); the other focuses more on *easy* samples to reduce the impact brought by noisy labels (e.g. self-paced learning [26,21]).

We study a different case: the test-time data distributions for different exits are divergent. Such distributions *could not be obtained in advance* due to the data-dependent inference procedure. The increased number of classifiers further raises challenges for designing the sample-weight mapping function. With our specially designed optimization objective and the meta-learning algorithm, we effectively produce proper sample weights for training the multi-exit model.

Meta-learning in sample weighting. Due to its remarkable progress, metalearning [42,8] has been extensively studied in sample weighting [37,38,50]. Existing approaches mainly focus on tackling class imbalance or corrupted label problems. In contrast, our goal is to improve the *inference efficiency* of early-exiting networks in a more general setting, without any assumption on the dataset.

3 Method

In this section, we first introduce the preliminaries of dynamic early-exiting networks and their conventional training strategies. Then the sample weighting mechanism and our meta-learning algorithm will be presented.

3.1 Preliminaries

Multi-exit networks. A typical approach to setting up a K-exit network is attaching K-1 intermediate classifiers at varying depths of a deep model [17,46]. For an input sample **x**, the prediction from the k-th exit can be written as

$$\hat{y}^{(k)} = \arg\max_{c} p_{c}^{(k)} = \arg\max_{c} f_{c}^{(k)}(\mathbf{x}; \boldsymbol{\Theta}_{f}^{(k)}), k = 1, 2, \cdots, K,$$
(1)

where $p_c^{(k)}$ is the k-th classifier's output probability for class c, and $f^{(k)}(\cdot; \Theta_f^{(k)})$ represents the k-th sub-network with parameter $\Theta_f^{(k)}$. Note that the parameters of different sub-networks are shared in a nested way except the classifiers. We denote the whole classification model as f and its parameters as Θ_f .



Fig. 2: Our training pipeline in iteration t. It consists of 3 steps: 1) the backbone network f^t is *pseudo-updated* to \hat{f}^t using training samples; 2) a meta objective is computed using \hat{f}^t on meta samples. This objective guides the update of the WPN g^t ; 3) f^t is updated using the new weights predicted by g^{t+1} .

Dynamic early exiting. Extensive efforts have been made to perform early exiting based on multi-exit models. A typical approach is terminating the forward propagation *once* the classification *confidence* $(\max_{c} p_{c}^{(k)})$ at a certain exit exceeds a given threshold (ϵ_{k}) . The final prediction is obtained by

$$\hat{y} = \hat{y}^{(k)}, \text{ if } \max_{c} p_{c}^{(k)} \ge \epsilon_{k}, \text{ and } \max_{c} p_{c}^{(j)} < \epsilon_{j}, \forall j \le k, k \le K - 1.$$
 (2)

The predictions $\hat{y}^{(k)}(k=1,2,\cdots,K)$ are obtained sequentially before satisfying the criterion in Eq. (2) or reaching the last exit. The threshold for exit-k (ϵ_k) can be decided on a validation set according to the computational budget.

Conventional training methods. A naive training strategy adopted by existing works [17,46] is directly minimizing a cumulative loss function:

$$\mathcal{L} = \sum_{k=1}^{K} \frac{1}{N} \sum_{i=1}^{N} l_{i}^{(k)} \triangleq \sum_{k=1}^{K} \frac{1}{N} \sum_{i=1}^{N} \operatorname{CE}(f^{(k)}(\mathbf{x}_{i}; \boldsymbol{\Theta}_{f}^{(k)}), y_{i}), \quad (3)$$

where CE is the cross-entropy loss, and N is the number of training samples.

3.2 Sample-weighting for early-exiting networks

In this subsection, we first formulate our sample weighting mechanism, and then introduce the proposed meta-learning objective. The optimization method is further presented. See Fig. 2 for an overview of the training pipeline.

Sample weighting with weight prediction network. We can observe from Eq. (2) that test samples are adaptively allocated to multiple exits according to their prediction confidence during inference. Sub-networks with varying depths are responsible for handling different subsets of samples. Therefore, it is suboptimal to set the same optimization objective for these exits as in Eq. (3). To this end, we propose to ameliorate the training objective by sample weighting:

$$\mathcal{L} = \sum_{k=1}^{K} \frac{1}{N} \sum_{i=1}^{N} w_i^{(k)} l_i^{(k)}.$$
(4)

Weight prediction network. Since we have no prior knowledge of where a specific sample exits, and the function mapping from input to weight is hard to establish,



Fig. 3: Updates of the backbone model f and the WPN g in detail. First, we compute the weighted classification loss (Eq. (5)), which guides the pseudo update on the backbone network f^t (Steps 1, 2, and 3). The meta objective (Eq. (6)) is then computed to update the WPN g^t (Steps 4 and 5).

we propose to automatically *learn* the weight $w_i^{(k)}$ from the input \mathbf{x}_i by a weight prediction network (WPN, denoted by g): $\mathbf{w}_i = [w_i^{(1)}, w_i^{(2)}, \cdots, w_i^{(K)}] = g(\mathbf{l}_i; \mathbf{\Theta}_g)$, where $\mathbf{l}_i = [l_i^{(1)}, l_i^{(2)}, \cdots, l_i^{(K)}]$ is the training loss for sample \mathbf{x}_i at K exits. The WPN g is established as an MLP with one hidden layer, and $\mathbf{\Theta}_g$ is the parameters of the WPN. We learn the backbone parameters $\mathbf{\Theta}_f$ and the WPN parameters $\mathbf{\Theta}_g$ in a meta-learning manner. Note that the WPN is only used for training, and no extra computation will be brought during the inference stage.

The meta-learning objective. We construct a novel optimization objective to bring the test-time adaptive behavior into training.

Weighted classification loss. With our weighting scheme (given Θ_g), the optimization objective of the backbone model parameter Θ_f can be written as

$$\boldsymbol{\Theta}_{f}^{\star}(\boldsymbol{\Theta}_{g}) = \arg\min_{\boldsymbol{\Theta}_{f}} \mathcal{L}_{\mathrm{tr}}(\boldsymbol{\Theta}_{f}, \boldsymbol{\Theta}_{g})$$
$$\triangleq \arg\min_{\boldsymbol{\Theta}_{f}} \sum_{k=1}^{K} \frac{1}{N} \sum_{i:\mathbf{x}_{i} \in \mathbf{X}_{\mathrm{tr}}} g^{(k)}(\boldsymbol{l}_{i}; \boldsymbol{\Theta}_{g}) \cdot l_{i}^{(k)}(\boldsymbol{\Theta}_{f}), \tag{5}$$

where \mathbf{X}_{tr} is the training set. Following [30], we scale the output of our WPN with $0 < \delta < 1$ to obtain a perturbation $\mathbf{\tilde{w}} \in [-\delta, \delta]^{B \times K}$, where *B* is the batch size. We further normalize the perturbation $\mathbf{\tilde{w}}$ to ensure the summation of its elements is 0. The final weight imposed on the loss is produced via $\mathbf{w} = 1 + \mathbf{\tilde{w}}$. *Meta data allocation.* The goal of our weighting mechanism is improving the model performance in the dynamic inference scenario. To construct the optimization objective for the WPN *g*, we first mimic the early exiting procedure on a meta set \mathbf{X}_{meta} , and obtain the meta samples exiting at different exits $\mathbf{X}_{meta}^{(k)}, k=1,2,\cdots,K$. Specifically, we define a budget controlling variable *q* to decide the number of samples that exit at each exit: $N_k = \frac{q^k}{\sum_{k=1}^K q^k} \times N_{\text{meta}}, q > 0$, where N_{meta} is the sample number in the meta set \mathbf{X}_{meta} . When q=1, the output numbers for different exits are equal, and q > 1 means more samples being output by deeper exits, together with more computation. In *training*, we generally tune the variable q on the validation set and fix it in the meta optimization procedure; in *testing*, we vary q from 0 to 2 to cover a wide range of computational budgets, and plot the curves of accuracy-computation in Sec. 4 (e.g. Fig. 5).

Given q, Fig. 3 (right) illustrates the procedure of obtaining $\mathbf{X}_{meta}^{(k)}$: each exit selects its most confident N_k samples which have not been output at previous classifiers. Note that for the last exit, $\mathbf{X}_{meta}^{(K)}$ contains all the unselected samples. *Meta objective for dynamic early exiting.* Instead of correctly recognizing all meta samples with diverse complexity, the meta objective for a specific exit is the classification loss on the samples that *actually exit* at this exit. The overall meta objective is obtained by aggregating the meta objectives from K exits:

$$\Theta_{g}^{\star} = \arg\min_{\Theta_{g}} \mathcal{L}_{\text{meta}}(\Theta_{f}^{\star}(\Theta_{g}))$$

$$\triangleq \arg\min_{\Theta_{g}} \sum_{k=1}^{K} \frac{1}{N_{k}} \sum_{j:\mathbf{x}_{j} \in \mathbf{X}_{\text{meta}}^{(k)}} l^{(k)}(\mathbf{x}_{j}, y_{j}; \Theta_{f}^{\star}(\Theta_{g})).$$
(6)

The optimization method. We jointly optimize the backbone parameter Θ_f (Eq. (5)) and the WPN parameter Θ_g (Eq. (6)) with an online strategy (see Fig. 2 for an overview of the optimization pipeline). Existing meta-learning methods for loss weighting typically require a standalone meta-set, which consists of class-balanced data or clean labels [37,38]. In contrast, we simply reuse the training data as our meta data. Precisely, at iteration t, we first split the input mini-batch into a training batch \mathbf{X}_{tr} and a meta batch \mathbf{X}_{meta} . The training batch \mathbf{X}_{tr} is used to construct the classification loss (Eq. (5)) and optimize the model parameters Θ_f . Next, the meta batch \mathbf{X}_{meta} is leveraged to compute the meta objective (Eq. (6)) and train the WPN parameters Θ_g .

Pseudo update of the backbone network is conducted using \mathbf{X}_{tr} :

$$\hat{\boldsymbol{\Theta}}_{f}^{t}(\boldsymbol{\Theta}_{g}) = \boldsymbol{\Theta}_{f}^{t} - \alpha \frac{\partial \mathcal{L}_{\mathrm{tr}}(\boldsymbol{\Theta}_{f}, \boldsymbol{\Theta}_{g})}{\partial \boldsymbol{\Theta}_{f}} \bigg|_{\boldsymbol{\Theta}_{f}^{t}}, \tag{7}$$

where α is the learning rate. Note that $\hat{\Theta}_{f}^{t}$ is a function of Θ_{g} , and this pseudo update is performed to construct the computational graph for later optimization of Θ_{g} . See Fig. 3 (gray dashed lines) for the gradient flow of this pseudo update. Update of the weight prediction network is performed using our meta objective calculated on \mathbf{X}_{meta} . Concretely, we mimic the early exiting procedure on \mathbf{X}_{meta} , and split it into K subsets without intersection $\mathbf{X}_{meta} = {\mathbf{X}_{meta}^{(1)} \cup \mathbf{X}_{meta}^{(2)} \cup \cdots \cup \mathbf{X}_{meta}^{(K)}}$, where $\mathbf{X}_{meta}^{(k)}$ contains the meta samples which should be output at exit-k according to the criterion in Eq. (2). See also Fig. 3 (right) for the procedure.

Receiving the partition of the meta-data based on the pseudo updated backbone network, the parameters of our weight prediction model can be updated: 8 Y. Han, Y. Pu et al.

$$\Theta_g^{t+1} = \Theta_g^t - \beta \frac{\partial \mathcal{L}_{\text{meta}}(\hat{\Theta}_f^t(\Theta_g))}{\partial \Theta_g} \bigg|_{\Theta_g^t}, \qquad (8)$$

where $\mathcal{L}_{\text{meta}}(\hat{\Theta}^t(\Theta_g))$ is the aggregation of the classification loss from each exit on its *allocated* meta data (Eq. (6)), and β is the learning rate. The gradient flow for updating Θ_g is illustrated in Fig. 3 (the golden dashed lines).

Algorithm 1 The meta-learning Algorithm

Require: Training data \mathcal{D} , batch size B, iteration T, interval I, budget controller q. **Ensure:** Backbone network parameters Θ_f . **for** t = 0 **to** T - 1 **do** { \mathbf{X}, \mathbf{y} } \leftarrow SampleMiniBatch(\mathcal{D}, B). Split { \mathbf{X}, \mathbf{y} } into { $\mathbf{X}_{tr}, \mathbf{X}_{meta}, \mathbf{y}_{tr}, \mathbf{y}_{meta}$ }. **if** $t \mod I = 0$ **then** Perform a pseudo update for $\hat{\Theta}_f^t$ by Eq. (7). Perform meta data allocation based on q. Update Θ_g^{t+1} by Eq. (8). **end if** Update Θ_f^{t+1} by Eq. (9). **end for**

Update of the backbone network is finally realized based on the updated Θ_a^{t+1} :

$$\Theta_f^{t+1} = \Theta_f^t - \alpha \frac{\partial \mathcal{L}_{tr}(\Theta_f, \Theta_g^{t+1})}{\partial \Theta_f} \bigg|_{\Theta_f^t}.$$
(9)

We summarize the learning algorithm in Algorithm 1. By mimicking the adaptive inference procedure, our novel objective in Eq. (6) encourages each exit to correctly classify the samples *which are most probably allocated to it* in the early exiting scenario (blue text in Eq. (6)). For example, early exits at shallow layers may focus more on those *easy* samples (see also our visualization results in Sec. 4, Fig. 8). More empirical analysis is presented in Sec. 4.

4 Experiments

In this section, we first conduct ablation studies to validate the design choices made in our weighting mechanism (Sec. 4.1). The main results on CIFAR [24] and ImageNet [5] are then presented in Sec. 4.2 and Sec. 4.3 respectively. Finally, we evaluate our method on the long-tailed CIFAR [4] in a class imbalance setting (Sec. 4.4). A network's performance is measured in terms of the trade-off between the accuracy and the Mul-Adds (multiply-add operations). We apply our training algorithm to two representative dynamic early-exiting architectures, i.e. multiscale dense network (MSDNet) [17] and resolution adaptive network (RANet) [46]. The experimental setup is provided in the supplementary material.



Fig. 4: Ablation studies on CIFAR-100. Left: our meta-learning based approach v.s. heuristic weighting mechanisms. Middle: the effectiveness of the proposed meta objective. Right: comparison of different WPN inputs.

4.1 Ablation studies

We perform ablation studies with a 5-exit MSDNet on CIFAR-100 to validate the effectiveness of different settings and variants in our method.

Meta-learning algorithm. We first verify the *necessity* of our meta-learning algorithm by comparing it with three variants: 1) the first replaces the weighting scheme with a vanilla selection scheme: the data allocation (Fig. 3) is applied directly to the *training* data, and the loss of each exit is only calculated on its *allocated* samples; 2) the second hand-designs a weighting mechanism with *fixed* weights increasing/decreasing with the exit index; 3) the third uses a *frozen* convergent WPN to weight the samples. For the second variant, the weights are set from 0.6 to 1.4 (or inverse) with uniform steps. Evaluation results on CIFAR-100 are shown in Fig. 4 (left). Several observations can be made:

- The vanilla sample selection strategy results in a drastic drop in accuracy, suggesting the necessity of our sample *weighting* mechanism;
- The gain of hand-crafted weight values is limited, indicating the advantage of our *learning*-based weighting approach;
- Interestingly, our joint optimization outperforms weighting with a frozen WPN. This suggests that it is essential to *dynamically* adjust the weights imposed on the loss at different training stages.

Meta-learning objective. Our designed meta loss in Eq. (6) encourages every exit to correctly recognize the data *subset* that is most likely handled by the exit in dynamic inference. To clarify the effectiveness of this meta objective, we keep the learning procedure unchanged and substitute Eq. (6) with the classification loss on the whole meta set for each exit. The results are shown in Fig. 4 (middle). While this variant (line Vanilla Objective) outperforms the baseline when the computation is relatively low, our objective achieves higher accuracy when larger computational budget is available. This suggests that emphasizing *hard* samples for deep classifiers is crucial to improving their performance.

The input of WPN. In addition to the classification loss, the confidence value can also be leveraged to produce the weight perturbations due to its role in making early exiting decisions. We test three types of input: 1) loss only; 2) confidence



Fig. 5: Dynamic inference results on CIFAR-10 (left) and CIFAR-100 (right)

Table 1: Anytime prediction results of a 7-exit MSDNet on CIFAR-100

Exit index		1	2	3	4	5	6	7
Params (×10 ⁶)		0.30	0.65	1.11	1.73	2.38	3.05	4.00
Mul-Add $(\times 10^6)$		6.86	14.35	27.29	48.45	76.43	108.90	137.30
Accuracy	MSDNet [17]	61.07	64.55	67.00	69.97	72.55	74.01	74.50
	Ours	$62.47 \ (\uparrow 1.30)$	$66.32 \ (\uparrow 1.77)$	$68.10 \ (\uparrow 1.10)$	$71.29 \ (\uparrow 1.32)$	$73.21 \ (\uparrow 0.66)$	$74.87 (\uparrow 0.86)$	$75.81 (\uparrow 1.31)$
	IMTA [29]	60.29	64.86	69.09	72.71	74.47	75.60	75.19
	Ours + IMTA [29]	62.26 (†1.97)	$67.18 (\uparrow 2.32)$	$70.53 (\uparrow 1.44)$	$73.10 (\uparrow 0.39)$	$74.80 (\uparrow 0.33)$	$76.05 (\uparrow 0.45)$	$76.31 \ (\uparrow 1.12)$

only; 3) the concatenation of loss and confidence. The accuracy-computation curves are shown in Fig. 4 (right). It can be found that the adoption of loss is essential for our method, and the inclusion of confidence could be harmful. We hypothesize that by reflecting the information of *both* network prediction and ground truth, the loss serves as a better candidate for WPN input. The ablation study of the WPN design is presented in the supplementary material.

4.2 CIFAR results

In this subsection, we report the results on CIFAR in both dynamic inference and anytime prediction settings following previous works [17,29,46]. We first apply our approach to MSDNet [17], and then compare it with the training techniques in [29]. The method is further validated on the RANet architecture [46].

Dynamic inference results. We apply our training strategy on MSDNet with 5 and 7 exits and compare with three groups of competitive baseline methods: classic networks (ResNet[13], DenseNet [18]), pruning-based approaches (Sparse Structure Selection (SSS) [19], Transformable Architecture Search (TAS) [6]), and dynamic networks (Shallow-Deep Networks (SDN) [23], Dynamic Convolutions (DynConv) [43], and Spatially Adaptive Feature Refinement (SAR) [12]).

In the dynamic inference scenario, we present the accuracy-computation (measured by Mul-Adds) curve in Fig. 5 (left: CIFAR-10, right: CIFAR-100). From the results, we can observe that the proposed weighting method consistently improves the performance of MSDNet at various computational budgets.



Fig. 6: **Dynamic inference results on CIFAR-100.** Left: comparison with IMTA [29]. Right: results with two different sized RANets.

Table 2: Anytime prediction results of a 6-exit RANet on CIFAR-100

Exit index		1	2	3	4	5	6
Params $(\times 10^6)$		0.36	0.90	1.30	1.80	2.19	2.62
Mul-Add ($\times 10^6$)		8.37	21.79	32.88	41.57	53.28	58.99
Accuracy	RANet [46]	63.41	67.36	69.62	70.21	71.00	71.78
	Ours	$65.33~(\uparrow 1.92)$	$68.69 (\uparrow 1.30)$	70.36 (†0.74)	$\textbf{70.80}~(\uparrow 0.59)$	$72.57 (\uparrow 1.57)$	$\textbf{72.45}~(\uparrow 0.67)$

When applied to the 5-exit model, our weighting mechanism obtains significant boosts with the CIFAR-100 Top-1 accuracy increased by about 1.3% when evoking around 15M Mul-Adds. For the 7-exit MSDNet, our model only uses half of the original budget to achieve ~94.0% Top-1 accuracy on CIFAR-10.

Anytime prediction results. We also report the accuracy of each exit on the whole test set in the anytime prediction setting. From the results in Tab. 1 (a 7-exit MSDNet [17]) and Tab. 2 (a 6-exit RANet [46]), we surprisingly find that although our meta objective encourages each exit to focus on only a *subset* of samples, the performance on the *whole* test set is consistently increased by a large margin. This phenomenon could bring some insights into the optimization of deep networks: 1) for the shallow exits, emphasizing a small subset of *easy* samples with *high confidence* benefits their generalization performance on the *whole* dataset [44]; 2) for the deeper exits, our objective forces them to focus on the *hard* samples, which cannot be confidently predicted by previous classifiers. Such a "challenging" goal could further improve their capability of approximating more complex classification boundaries. This coincides with the observation in our ablation studies (Fig. 4 middle): encouraging each exit to correctly classify a *well-selected* subset of training samples is preferable.

Comparison and compatibility with IMTA. Improved techniques for training adaptive networks (IMTA) [29] are proposed to stabilize training and facilitate the collaboration among exits. These techniques are developed for the *optimization procedure* and ignore the adaptive inference behavior. In contrast,



Fig. 7: **ImageNet results.** Top left: accuracy-computation curves. Others: accuracy-latency curves on different hardware platforms.

our meta-learning based method improves the *objective optimization* and takes the inference behavior into account. We empirically compare our approach to IMTA [29], and further combine them to evaluate the compatibility.

The results are presented in Fig. 6 (right, for dynamic inference) and Tab. 1 (row 6 & 7, for anytime prediction). These results validate that 1) our weighting strategy achieves higher performance, especially at low computational budges; 2) combining our method with IMTA further improves the performance. In particular, note that the accuracy of deep exits can be boosted by a large margin.

Results on RANet. To evaluate the generality of our method, we conduct experiments on another representative multi-exit structure, RANet [46]. The anytime prediction performance of a 6-exit RANet is shown in Tab. 2, and the dynamic inference results of two different-sized RANets are illustrated in Fig. 6 (right). These results suggest that the proposed approach consistently improves the trade-off between accuracy and computational cost of RANet. Notably, when the Mul-Add is around 15M, our model improves the Top-1 accuracy significantly (~1.4%). The strong performance on MSDNet and RANet indicates that our weighting algorithm is sufficiently general and effective.



Fig. 8: Visualization results of the weight-loss relationship on ImageNet

4.3 ImageNet results

Next, we evaluate our method on the large-scale ImageNet dataset [5]. We test three different sized MSDNets following [17] and compare them with competitive baselines, including pruning-based approaches (the aforementioned SSS [19] and TAS [6], and Filter Pruning via Geometric Median (Geom) [14]) and dynamic networks (SkipNet [45], Batch-Shaping (BAS) [1], Channel Selection [15] and DynConv [43]). We report the results in the dynamic inference setting here. Anytime prediction results can be found in the supplementary material.

Accuracy-computation result. We present the accuracy-computation results on ImageNet in Fig. 7 (top left). The plot shows that the weighting mechanism consistently improves the accuracy-computation trade-off of MSDNets. Even though some competing models surpass the baseline MSDNet trained without sample weighting, our method successfully outperforms these competitors.

Accuracy-latency result. We benchmark the efficiency of our models on three types of computing devices: Intel i5-1038NG7 mobile CPU, Nvidia Jetson TX2 (a resource-constrained embedded AI computing device) and TITAN Xp GPU. The testing batch sizes are set as 64 for the first two devices and 512 for the last. The testing images have resolutions of 224×224 pixels. The accuracy-latency curves plotted in Fig. 7 demonstrate the significant improvement of our method across all computing platforms. For instance, our model only takes 70% of the original MSDNet's computation time to achieve 75.4% accuracy ($1.4 \times$ speed-up). Visualization. We visualize the weights predicted by WPN toward the end of training in Fig. 8⁴. The left plot presents the weight-loss relationship at exit-1 and exit-3. The right plot shows the weights of the samples selected/not selected by exit-1 in meta-data allocation (Sec. 3.2). Several observations can be made: 1) Since more samples are allocated to early exits, the weights at exit-1 are generally larger than those at exit-3. This result suggests that the proposed method successfully learns the relationship between exits; 2) Both exits tend to emphasize the samples with a smaller loss. This coincides with the observation

⁴ We set q = 0.5 in training, and therefore the proportion of output samples at 5 exits follows an exponential distribution of [0.52, 0.26, 0.13, 0.06, 0.03].



Fig. 9: Results on long-tailed CIFAR-100 with different imbalance factors

in [51] which explains the importance of the well-classified (*easy*) samples in the optimization of deep models; 3) From weights at exit-1, we can observe that the samples with high prediction confidence (which would be selected by exit-1 in meta-data allocation) generally have weights larger than samples with lower confidence (which would be allocated to deeper classifiers). This suggests that the proposed method learns to recognize the samples' exit at different classifiers.

4.4 Class imbalance results

The proposed method is finally evaluated in the class imbalance setting. On the long-tailed CIFAR-100 [4] dataset, we test with four imbalance factors (200, 100, 50, and 20), where the imbalance factor is defined as the number of training samples in the largest class divided by the smallest. The budgeted batch classification performance illustrated in Fig. 9 shows that our weighting method consistently outperforms the conventional training scheme by a large margin.

5 Conclusion

In this paper, we propose a meta-learning based sample weighting mechanism for training dynamic early-exiting networks. Our approach aims to bring the test-time adaptive behavior into the training phase by sample weighting. We weight the losses of different samples at each exit by a weight prediction network. This network is jointly optimized with the backbone network guided by a novel meta-learning objective. The proposed weighting scheme can consistently boost the performance of multi-exit models in both anytime prediction and budgeted classification settings. Experiment results validate the effectiveness of our method on the long-tailed image classification task.

Acknowledgement. This work is supported in part by National Key R&D Program of China (2020AAA0105200), the National Natural Science Foundation of China under Grants 62022048, Guoqiang Institute of Tsinghua University and Beijing Academy of Artificial Intelligence. We also appreciate the generous donation of computing resources by High-Flyer AI.

References

- 1. Bejnordi, B.E., Blankevoort, T., Welling, M.: Batch-shaping for learning conditional channel gated networks. In: ICLR (2020)
- 2. Bolukbasi, T., Wang, J., Dekel, O., Saligrama, V.: Adaptive neural networks for efficient inference. In: ICML (2017)
- Chawla, N.V., Bowyer, K.W., Hall, L.O., Kegelmeyer, W.P.: Smote: synthetic minority over-sampling technique. JAIR (2002)
- 4. Cui, Y., Jia, M., Lin, T.Y., Song, Y., Belongie, S.: Class-balanced loss based on effective number of samples. In: CVPR (2019)
- 5. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR (2009)
- Dong, X., Yang, Y.: Network pruning via transformable architecture search. NeurIPS (2019)
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. In: ICLR (2021)
- 8. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: ICML (2017)
- 9. Freund, Y., Schapire, R.E.: A decision-theoretic generalization of on-line learning and an application to boosting. JCSS (1997)
- 10. Han, S., Mao, H., Dally, W.J.: Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding (2016)
- Han, Y., Huang, G., Song, S., Yang, L., Wang, H., Wang, Y.: Dynamic neural networks: A survey. TPAMI (2021)
- 12. Han, Y., Huang, G., Song, S., Yang, L., Zhang, Y., Jiang, H.: Spatially adaptive feature refinement for efficient inference. TIP (2021)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
- 14. He, Y., Liu, P., Wang, Z., Hu, Z., Yang, Y.: Filter pruning via geometric median for deep convolutional neural networks acceleration. In: CVPR (2019)
- Herrmann, C., Bowen, R.S., Zabih, R.: Channel selection using gumbel softmax. In: ECCV (2020)
- Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: Mobilenets: Efficient convolutional neural networks for mobile vision applications. arXiv preprint arXiv:1704.04861 (2017)
- 17. Huang, G., Chen, D., Li, T., Wu, F., van der Maaten, L., Weinberger, K.: Multiscale dense networks for resource efficient image classification. In: ICLR (2018)
- Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: CVPR (2017)
- Huang, Z., Wang, N.: Data-driven sparse structure selection for deep neural networks. In: ECCV (2018)
- Hubara, I., Courbariaux, M., Soudry, D., El-Yaniv, R., Bengio, Y.: Binarized neural networks. In: NeurIPS (2016)
- Jiang, L., Meng, D., Mitamura, T., Hauptmann, A.G.: Easy samples first: Selfpaced reranking for zero-example multimedia search. In: ACM MM (2014)
- 22. Johnson, J.M., Khoshgoftaar, T.M.: Survey on deep learning with class imbalance. Journal of Big Data (2019)

- 16 Y. Han, Y. Pu et al.
- Kaya, Y., Hong, S., Dumitras, T.: Shallow-deep networks: Understanding and mitigating network overthinking. In: ICML (2019)
- Krizhevsky, A.: Learning multiple layers of features from tiny images. Tech. rep. (2009)
- Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: NeurIPS (2012)
- Kumar, M., Packer, B., Koller, D.: Self-paced learning for latent variable models. In: NeurIPS (2010)
- 27. LeCun, Y., Denker, J.S., Solla, S.A.: Optimal brain damage. In: NeurIPS (1990)
- Li, H., Kadav, A., Durdanovic, I., Samet, H., Graf, H.P.: Pruning filters for efficient convnets. In: ICLR (2017)
- 29. Li, H., Zhang, H., Qi, X., Yang, R., Huang, G.: Improved techniques for training adaptive deep networks. In: ICCV (2019)
- Li, S., Ma, W., Zhang, J., Liu, C.H., Liang, J., Wang, G.: Meta-reweighted regularization for unsupervised domain adaptation. TKDE (2021)
- 31. Lin, J., Rao, Y., Lu, J., Zhou, J.: Runtime neural pruning. In: NeurIPS (2017)
- Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: ICCV (2017)
- 33. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: ICCV (2021)
- Liu, Z., Li, J., Shen, Z., Huang, G., Yan, S., Zhang, C.: Learning efficient convolutional networks through network slimming. In: ICCV (2017)
- Malisiewicz, T., Gupta, A., Efros, A.A.: Ensemble of exemplar-syms for object detection and beyond. In: ICCV (2011)
- 36. Panda, P., Sengupta, A., Roy, K.: Conditional deep learning for energy-efficient and enhanced pattern recognition. In: DATE (2016)
- 37. Ren, M., Zeng, W., Yang, B., Urtasun, R.: Learning to reweight examples for robust deep learning. In: ICML (2018)
- 38. Shu, J., Xie, Q., Yi, L., Zhao, Q., Zhou, S., Xu, Z., Meng, D.: Meta-weight-net: Learning an explicit mapping for sample weighting. In: NeurIPS (2019)
- Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: ICLR (2015)
- 40. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: CVPR (2015)
- 41. Teerapittayanon, S., McDanel, B., Kung, H.T.: Branchynet: Fast inference via early exiting from deep neural networks. In: ICPR (2016)
- 42. THospedales, T., Antoniou, A., Micaelli, P., Storkey, A.: Meta-learning in neural networks: A survey. TPAMI (2021)
- 43. Verelst, T., Tuytelaars, T.: Dynamic convolutions: Exploiting Spatial Sparsity for Faster Inference. In: CVPR (2020)
- 44. Wang, X., Chen, Y., Zhu, W.: A survey on curriculum learning. TPAMI (2021)
- 45. Wang, X., Yu, F., Dou, Z.Y., Darrell, T., Gonzalez, J.E.: Skipnet: Learning dynamic routing in convolutional networks. In: ECCV (2018)
- 46. Yang, L., Han, Y., Chen, X., Song, S., Dai, J., Huang, G.: Resolution Adaptive Networks for Efficient Inference. In: CVPR (2020)
- Yang, L., Jiang, H., Cai, R., Wang, Y., Song, S., Huang, G., Tian, Q.: Condensenet v2: Sparse feature reactivation for deep networks. In: CVPR (2021)
- Zadrozny, B.: Learning and evaluating classifiers under sample selection bias. In: ICML (2004)
- 49. Zhang, X., Zhou, X., Lin, M., Sun, J.: Shufflenet: An extremely efficient convolutional neural network for mobile devices. In: CVPR (2018)

- 50. Zhang, Z., Pfister, T.: Learning fast sample re-weighting without reward data. In: ICCV (2021)
- 51. Zhao, G., Yang, W., Ren, X., Li, L., Sun, X.: Well-classified examples are underestimated in classification with deep neural networks. In: AAAI (2022)